




National Aeronautics and  
Space Administration

Jet Propulsion Laboratory  
California Institute of Technology  
Pasadena, California



National Aeronautics and  
Space Administration

Jet Propulsion Laboratory  
California Institute of Technology  
Pasadena, California

# Using S3 in Apache Science Data Analytics Platform

Presented By: Frank Greguska (JPL)

Jet Propulsion Laboratory  
California Institute of Technology  
4800 Oak Grove Drive, Pasadena, CA 91109-8099, U.S.A.



- What is Apache Science Data Analytics (SDAP)?
- SDAP Analytics and Architecture
- Data Storage Comparison
- Benchmarks

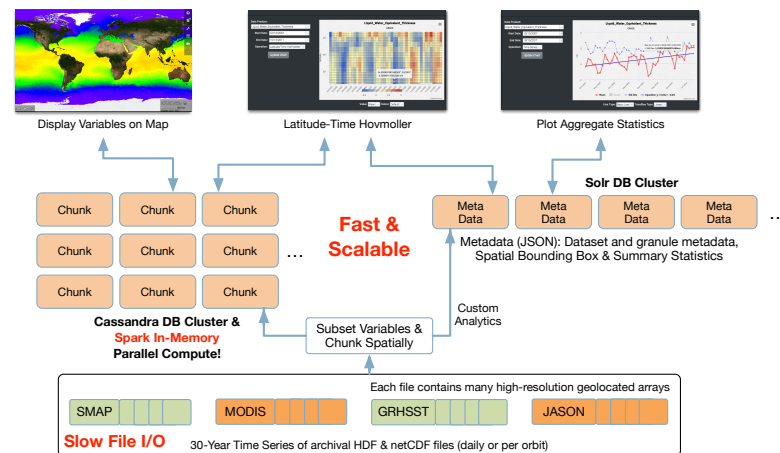
# Apache Science Data Analytics Platform (SDAP)

- **OceanWorks** is to establish an **Integrated Data Analytics Center** at the NASA Physical Oceanography Distributed Active Archive Center (PO.DAAC) for Big Ocean Science
- Focuses on technology integration, advancement and maturity
- Collaboration between JPL, Center for Atmospheric Prediction Studies (COAPS) at Florida State University (FSU), National Center for Atmospheric Research (NCAR), and George Mason University (GMU)
- Bringing together PO.DAAC-related big data technologies
  - Big data analytic platform
  - Anomaly detection and ocean science
  - Distributed in situ to satellite matchup
  - Dynamic datasets ranking and recommendations
  - Sub-second data search solution and metadata translation and services aggregation
  - Quality-screened data subsetting
- All code open-sourced as Apache Science Data Analytics Platform (SDAP)



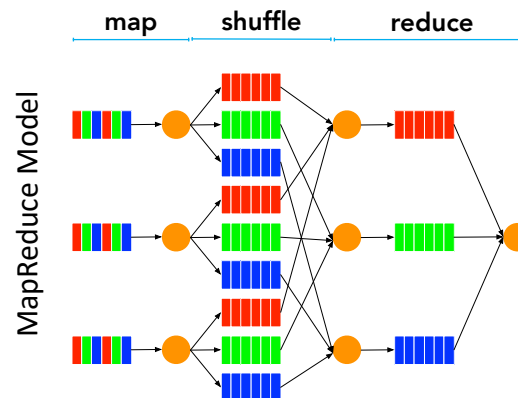
# SDAP Cloud Analytics: NEXUS

- **NEXUS** is a data-intensive analysis solution using a new approach for handling science data to enable large-scale data analysis
  - Streaming architecture for horizontal scale data ingestion
  - Scales horizontally to handle massive amount of data in parallel
  - Provides high-performance geospatial and indexed search solution
  - Provides tiled data storage architecture to eliminate file I/O overhead
  - A growing collection of science analysis webservice



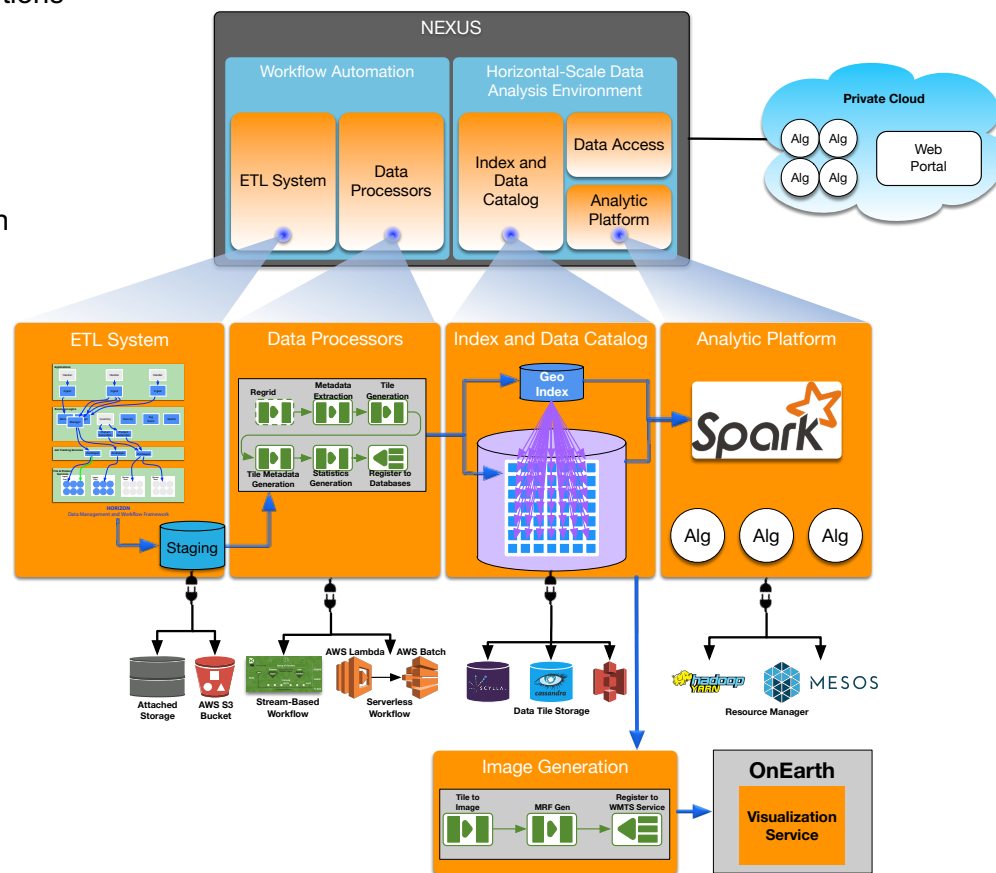
NEXUS' Two-Database Architecture

- **MapReduce**: A programming model for expressing distributed computations on massive amount of data and an execution framework for large-scale data processing on clusters of commodity servers. - J. Lin and C. Dyer, "*Data-Intensive Text Processing with MapReduce*"
  - **Map**: splits processing across cluster of machines in parallel, each is responsible for a record of data
  - **Reduce**: combines the results from Map processes





# NEXUS' Pluggable Architecture

- NEXUS supports public/private Cloud and local cluster deployments
- It has a growing set of algorithms – Time Series | Latitude/Time Hovmöller| Longitude/Time Hovmöller| Latitude/Longitude Time Average | Area Averaged Time Series | Time Averaged Map | Climatological Map | Correlation Map | Daily Difference Average
- It offers several container-based deployment options
  - Local on-premise cluster
  - Private Cloud
  - Amazon Web Service
- Automate Data Ingestion with Image Generation
  - Cluster based
  - Serverless (Amazon Lambda and Batch)
- Data Store Options**
  - Apache Cassandra**
  - Amazon Simple Storage Service (S3)**
- Resource Management Options
  - Apache YARN
  - Apache MESOS
- Analytic Engine Options
  - Custom Apache Spark Cluster
  - Amazon Elastic MapReduce (EMR)
  - Amazon Athena (work-in-progress)



# Cassandra Vs. S3

- Apache Cassandra
  - NoSQL database cluster where data is stored in Amazon's Elastic Block Store (EBS)
  - EBS is a block storage (which means the type of file system must be configured) that is attached to an Amazon EC2 instance
  - EBS is only available in the same region as the EC2 instance it is attached to; it is basically an external hard drive
  - Faster than S3 but not as scalable
- Amazon S3
  - A standalone object store (not a filesystem; it has a flat structure)
  - Highly available and redundant. The chance of data loss is basically 0%
  - Much cheaper than EBS/Cassandra implementation
  - Slower than EBS but way more scalable!

	Amazon EBS 	Amazon S3 
Paradigm	File system	Object store
Performance	Very fast	Fast
Redundancy	Across multiple servers in an Availability Zone	Across multiple facilities and on multiple devices within each facility
Security	EBS Encryption – Data volumes and Snapshots	Server Side Encryption
Access from the Internet?	No (1)	Yes (2)
Typical use case	It is a disk drive	Write once, read many

(1) Accessible from the Internet if mounted to server and set up as FTP, etc.

(2) Only with proper credentials, unless ACLs are world-readable

# Cassandra Vs. S3

## Ingestion

- Ingestion process needs to save tiles to the correct Data Store
- Easily choose Cassandra or S3 with configuration options

```
# Tile writer configuration
ningester:
  tile_writer:
    data_store: s3Store
    metadata_store: solrStore
---
spring:
  data:
    solr:
      host: http://127.0.0.1:8983/solr/

datasource:
  s3Store:
    region: us-west-2
    bucketName: sdap-tiles
  solrStore:
    collection: nexustiles
```

*Example connection settings yaml file*

## Analysis

- Analysis needs to read tiles from the correct data store
- Easily choose Cassandra or S3 with configuration options

```
[cassandra]
host=127.0.0.1
keyspace=nexustiles
local_datacenter=datacenter1
protocol_version=3

[s3]
bucket=sdap-tiles
region=us-west-2

[solr]
host=localhost:8983
core=nexustiles

[datastore]
store=s3
```

*Example connection settings ini file*



# Cassandra Vs. S3


## Which one is better?






# Benchmark Results (Ingestion)

- AVHRR SST:






- MODIS AOD:



*Benchmarks provided by David Silvan*

# Benchmark Results (API Endpoints)



- AVHRR SST:



Benchmarks provided by David Silvan





# Benchmark Results (API Endpoints)

- AVHRR SST:



Benchmarks provided by David Silvan

- MODIS AOD (2002-2016):



Benchmarks provided by David Silvan



# Conclusion

- OceanWorks is open sourced as Apache Science Data Analytics Platform (SDAP)
  - <http://sdap.apache.org>
- SDAP NEXUS enable analysis of data stored in Cassandra or S3
- Configuring SDAP NEXUS to use Cassandra or S3 is easy
- Cassandra ranges between 1.5x – 15x faster than S3
- Cassandra on EBS storage is more expensive than S3
- Amazon S3 can be useful for reducing cost of compute cluster at the expense of performance
- Cassandra should be used when performance is of major concern